Product use case

Veeam Kasten Use Case

Kubernetes-native Backup and Restore

New Considerations for Kubernetes Container Technology

Virtualization and cloud technologies have triggered applications and services to outgrow the traditional models as a reference design. Kubernetes, a virtualized container platform, is an example of a new technology that requires new tools and support systems, including backup and restore. The microservices framework is a new model that characterizes the structure of applications that are intended for cloud environments.

Traditional tools and support systems are not adequate to automatically capture an application built on the microservices architecture. Therefore, restoring an application with traditional tools is not practical.

Service Restoration is Key

Backup and restore refers to a process or solution that generates a copy of an application in production and stores it to a separate infrastructure resource. Ideally, the backup location is isolated from the production platform to create independence, so that a production system infrastructure failure does not concurrently impact the stored copy, or backup. This backup can then be used to return a service to operation from the point that the snapshot was created.

Restore Microservices to Live Production

As cloud-native adoption, containerization, and microservices deployment continues to surge, Kubernetes is emerging as the most rapidly growing infrastructure platform.

Restoring a microservices application to its operational state necessitates abstracting the system from its underlying infrastructure while considering data services that exist outside of the cluster. Consequently, the need for Kubernetes native backup solutions becomes imperative.

Key Veeam Kasten Capabilities

Application-specific Protection

Veeam Kasten automatically discovers all Kubernetes application components and provides the option to back up each of the components according to specified parameters. This option not only helps to secure data used by applications, but helps to optimize backup storage usage too.

Policy-based Backup Automation and Monitoring

Policies automate data management workflows at scale. To achieve this, policies combine actions set by the administrator (e.g., take a snapshot), at a frequency or schedule for that action. Users can define a namespace, or a label-based selection criteria for the resources set for backup, to optimize protection.

Rapid Service Restoration

One of the main goals of an effective Kubernetes backup solution is the non-invasive capture of containerized applications, coupled with the ability to restore it quickly and effectively. Veeam Kasten restores applications through a variety of options, including selective restore, to maximize the use of stored data for integration with normal restore and disaster recovery (DR) operations.

© 2024 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Product use case

Kubernetes and Traditional Workloads are Based on Different Architectures

Traditional Client-Server Architecture	VS	Kubernetes Microservices Architecture
A three-tiered framework consisting of presentation, application, and database layers	VS	A cloud application is a collection of decoupled services and data
Updates require an incremental production cycle	VS	An application has no set data interfaces
The entire application is deployed as a single module and ignores persistent data stored outside the volume	VS	Each microservice can be written in its own language and the application is captured collectively
Service restoration is dependent on the complete volume and may take days	VS	Returning a service to production on a secure platform can take just minutes

Backup and Restore Use Case Scenarios

Application Misconfiguration

When applications are misconfigured, the results may not be immediately detected upon the release of new code or a software update. Misconfigurations are detected only after an expected service result is not in alignment with a service or application specification.

Accidental Data Loss

Accidental data loss can occur throughout development and production cycles. Even with tight development controls built into software production, dependent components such as data stores can become corrupted or accidentally deleted.

Standby Clusters

Running redundant operations is a method to protect against failure that may occur to systems in production. In Kubernetes applications, these are called "standby clusters," and the tools that enable software clusters to be copied, migrated, or otherwise placed into service are important. A complete backup solution will have the capability to copy, restore, and move clusters, as well have an option to source a combination of offline and offsite protected data for service restoration.

Compliance

Backup and restore capabilities as well as data protection practices are governed by a combination of internal policies and external regulatory control. Because the new microservices architectures embraced by the Kubernetes platform are not automatically compliant with traditional volume backup tools, only Kubernetes backup solutions can address operations in compliance with most policies.

© 2024 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Product use case

Backup and Restore Specification Sheet

Backup and Restore	Specifications
Kubernetes distribution(s)	• Amazon Elastic Kubernetes Service (EKS) • Microsoft Azure Kubernetes Service (AKS) • Digital Ocean • Kubernetes • Nutanix • Oracle Container Engine (OKE) • SUSE Rancher • Amazon EKS-Anywhere • Google Kubernetes) • HPE Ezmeral • K3s • Red Hat OpenShift • VMware Tanzu
Marketplaces	• AWS Marketplace, Azure Marketplace, Red Hat Marketplace, SUSE Rancher Marketplace, VMware vSphere with Tanzu
Kubernetes extensions	• HELM (installation, chart signing), Kanister, Kopia, Kubestr, Navig8
Data sources	• Amazon RDS, Cassandra, Elasticsearch, Kanister, Kafka, K8ssandra, MongoDB, MySQL, PostgreSQL, SQL Server
Data formats	• JSON, XML, YAML,
On-premises storage systems	• Dell-EMC, HPE, IBM, Lenovo, NetApp, Pure Storage
Cloud storage	• Amazon EBS -EFS -S3, Azure Disk Storage, Ceph, Cisco HyperFlex, CSI, Dell EMC, Google Cloud Storage, Hitachi, HPE, Infinidat, Lenovo, MinIO, Net App, Oracle OCI, Pure Storage, Zadara
Security	• FIPS 140-3, RBAC, OIDC, Token Auth, IAM, ransomware immutability
Security integrations	• Kyverno, Open Policy Agent, Red Hat, Hashi Corp Vault, AWS
Public cloud platform	• AWS, Google Cloud, Azure, Digital Ocean
Monitoring integrations	• Azure Sentinel, DataDog, Prometheus
Observability	• Out-of-the-box dashboards, reports, metrics
Number of clusters	• 5 (free) / 500 (trial) / Unlimited (paid)
Backup range settings	• Continuous to monthly, specified hours, day, week, month

© 2024 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Product use case

Veeam Kasten *for Kubernetes* Data Protection Platform

Veeam Kasten is a Kubernetes-native data protection platform that provides enterprise operations teams with an easy to use, scalable, and secure system for backup and restore, disaster recovery (DR), and application mobility of Kubernetes applications. With Veeam, teams achieve Kubernetes-native resilience against ransomware attacks. Kasten offers an application-centric approach and deep integrations with relational and NoSQL databases, Kubernetes distributions and APIs, and cloud platforms. This provides teams with the freedom to choose any infrastructure they want while achieving maximum operational simplicity. Policy-driven and extensible, Kasten also includes features such as full-spectrum consistency, database integrations, automatic application discovery, multi-cloud mobility, and a powerful web-based user interface.

> "Throughout our evaluation, Enterprise Strategy Group determined that Kasten can deliver the data protection and recovery capabilities required by Kubernetes applications. We examined this by comparing how traditional solutions and Kasten address four major use cases: Backup and restore, DR, application mobility, and ransomware protection."

Dispelling the Myths of Kubernetes Data Protection
ESG 04 2024

Kubernetes Data Protection Use Cases: Summary

Veeam deeply understands Kubernetes and it's unique challenges with regard to backup, restore, DR, application mobility, and ransomware protection. Veeam Kasten helps enterprises successfully run applications on Kubernetes with confidence.

Veeam Kasten Use Cases

Backup & Restore

Protect your cloud native Kubernetes and VM applications, while preserving your business-critical data.

Disaster Recovery

Manage how backups are replicated off-site to meet business and regulatory requirements.

Application Mobility

Move applications between clouds and on-premises for test/dev, load balancing, data management, and upgrades.

Ransomware Protection

Protect your Kubernetes platform during cyberattacks to preserve business continuity.

For more information, visit Veeam.com or follow @Veeam on X.

© 2024 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.